

Naturalness of MSSM dark matter

María Eugenia Cabrera,^a J. Alberto Casas,^b Antonio Delgado,^c Sandra Robles,^{b,d} and Roberto Ruiz de Austri^e

^a*Instituto de Física, Universidade de São Paulo, São Paulo SP 05508-900, Brazil*

^b*Instituto de Física Teórica, IFT-UAM/CSIC, Universidad Autónoma de Madrid, Cantoblanco, 28049 Madrid, Spain*

^c*Department of Physics, University of Notre Dame, Notre Dame, IN 46556, USA*

^d*Departamento de Física Teórica, Universidad Autónoma de Madrid, Cantoblanco, 28049 Madrid, Spain*

^e*Instituto de Física Corpuscular, IFIC-UV/CSIC, Valencia, Spain*

E-mail: mcabrera@if.usp.br, alberto.casas@uam.es,
antonio.delgado@nd.edu, sandra.robles@uam.es, r Ruiz@ific.uv.es

ABSTRACT: There exists a vast literature examining the electroweak (EW) fine-tuning problem in supersymmetric scenarios, but little concerned with the dark matter (DM) one, which should be combined with the former. In this paper, we study this problem in an, as much as possible, exhaustive and rigorous way. We have considered the MSSM framework, assuming that the LSP is the lightest neutralino, χ_1^0 , and exploring the various possibilities for the mass and composition of χ_1^0 , as well as different mechanisms for annihilation of the DM particles in the early Universe (well-tempered neutralinos, funnels and co-annihilation scenarios). We also present a discussion about the statistical meaning of the fine-tuning and how it should be computed for the DM abundance, and combined with the EW fine-tuning. The results are very robust and model-independent and favour some scenarios (like the h-funnel when $M_{\chi_1^0}$ is not too close to $m_h/2$) with respect to others (such as the pure wino case). These features should be taken into account when one explores “natural SUSY” scenarios and their possible signatures at the LHC and in DM detection experiments.

KEYWORDS: Supersymmetry Phenomenology

Contents

1	Introduction	1
2	The measure of the fine-tuning	3
2.1	Assumptions behind the standard fine-tuning criterion	3
2.2	Examples	5
3	Well tempered bino-Higgsino	5
4	Well-tempered bino-wino(-Higgsino)	10
5	Funnels	14
6	Annihilation and Co-annihilation	18
7	Connection to the electroweak fine-tuning	20
8	Accommodating Higgsino DM in the MSSM	22
9	Conclusions	23

1 Introduction

In the minimal supersymmetric Standard Model (MSSM), there are several potential sources of fine-tuning. The most notorious one is the electroweak (EW) fine-tuning, which generically requires light gluino, light Higgsinos, (not so) light winos and, in many cases, light stops. This fine-tuning can be reasonably quantified by the “standard” measure [1, 2]:

$$\Delta_i^{(\text{EW})} = \frac{d \log v^2}{d \log \theta_i} ; \quad \Delta^{(\text{EW})} \equiv \max \left\{ \Delta_i^{(\text{EW})} \right\}, \quad (1.1)$$

where v^2 is the Higgs vacuum expectation value (VEV) and θ_i are the independent (initial) parameters of the model under consideration. Typically $\Delta^{(\text{EW})}$ is dominated by the gluino-mass parameter and its value is $\gtrsim \mathcal{O}(100)$ [3], corresponding to a fine-tuning at the level of $\lesssim 1\%$. There is a vast literature concerning this EW fine-tuning of the MSSM. An important fact is that $\tan \beta$ should be moderately large (say $\tan \beta \gtrsim 6$) in order to reproduce the experimental Higgs mass without the need of gigantic stop masses, which would imply a very severe fine-tuning.

Besides the EW fine-tuning, there is a potential fine-tuning related to the generation of the right amount of dark matter (DM). In some scenarios of supersymmetric dark matter, a delicate balance between a-priori-independent quantities is required, denoting a fine-tuned situation. Here, by contrast, the literature is much less extensive [4–7] and, furthermore, many important mechanisms of supersymmetric dark matter have never been considered from this point of view¹. The main goal of this paper is precisely to perform a rigorous study of the fine-tuning associated with the production of MSSM dark matter in all the interesting scenarios. Moreover, we will combine this fine-tuning with the EW one, to select the MSSM regions that are globally less fine-tuned.

¹For works studying the effect of DM constraints on the EW fine-tuning, see e.g. [8, 9].

We will focus on the case where the DM particle is a supersymmetric WIMP, namely the lightest state of the neutralino mass matrix,

$$\mathcal{M}_{\chi^0} = \begin{pmatrix} M_1 & 0 & -m_Z s_W c_\beta & m_Z s_W s_\beta \\ 0 & M_2 & m_Z c_W c_\beta & -m_Z c_W s_\beta \\ -m_Z s_W c_\beta & m_Z c_W c_\beta & 0 & -\mu \\ m_Z s_W s_\beta & -m_Z c_W s_\beta & -\mu & 0 \end{pmatrix}, \quad (1.2)$$

which is the “standard” situation. Of course, the lightest neutralino, χ_1^0 , must also be the lightest supersymmetric particle (LSP). In the previous equation, M_1 and M_2 are the (low-energy) bino and wino soft mass parameters, while μ is the mass parameter in the superpotential, which gives mass to Higgsinos. As usual, s_W (c_W) is the sin (cosine) of the weak angle and s_β (c_β) is the sin (cosine) of the β -angle, defined by the ratio of the two Higgs VEVs, $\tan \beta = \langle H_u \rangle / \langle H_d \rangle$. Generically, χ_1^0 is a combination of bino, wino and Higgsinos, though is usually dominated by one of these species. Certainly, the content of χ_1^0 in each species depends on the particular values of the four parameters that define \mathcal{M}_{χ^0} , i.e. $\{M_1, M_2, \mu, \tan \beta\}$.

The lightest neutralino is a perfect candidate for DM, but, to be successful, it must be produced in the early Universe in the right amount to reproduce the present DM relic density [10]

$$\Omega_{\text{DM}}^{(\text{obs})} h^2 = 0.119 \pm 0.012. \quad (1.3)$$

We will suppose, throughout this paper, that the neutralino relic density was produced in the “standard” thermal way, i.e. under the assumptions that neutralinos were produced thermally thanks to their interactions with other particles in the primordial plasma, and that they decoupled while the Universe was radiation-dominated. Then, their present relic density is given by [11]

$$\Omega_{\text{DM}} h^2 = \frac{8.7 \times 10^{-11} \text{ GeV}^{-2}}{\sqrt{g^*} \int_{x_f}^{\infty} \langle \sigma_{\text{ann}} v \rangle x^{-2} dx}, \quad (1.4)$$

where the g^* parameter accounts for the number of degrees of freedom at freeze-out, $x \equiv m/T$, i.e. temperature over mass, and the subscript f denotes the freeze-out time, $T_f \simeq m/20$. Besides, $\langle \sigma_{\text{ann}} v \rangle$ stands for the thermal-averaged annihilation cross section (times the velocity). Thus, in order to reproduce the observed relic density (1.3) the neutralinos must annihilate at early times with a suitable cross section.

From the naturalness point of view, an interesting case occurs when χ_1^0 is close to a pure state. Then, roughly speaking, $\sigma \propto m_{\chi_1^0}^{-2}$ and therefore, in order to reproduce (1.3), there is in principle no need of any fine-arrangement of the parameters in the \mathcal{M}_{χ^0} matrix; only a particular value of $m_{\chi_1^0}$, i.e. $\sim M_1$, M_2 or μ , depending on the character of χ_1^0 . Actually, the case of (close to) pure bino does not work, since its annihilation rate in the early Universe is typically too small for any value of M_1 , leading to an overproduction of dark matter, totally inconsistent with eq. (1.3). In contrast, the cases of (essentially) pure Higgsino or pure wino lead to the correct relic density if their masses are, respectively, $\mu \simeq 1 \text{ TeV}$ or $M_2 \simeq 3 \text{ TeV}$.²

Notice that both cases lead to a rather heavy supersymmetric spectrum, which has two problems. First of all, the expectations to discover supersymmetry at the LHC decrease (actually, for the wino-LSP they vanish). Second, the heavier the spectrum, the more fine-tuned the model with respect to the EW breaking. It is therefore of interest to consider mechanisms that allow for lighter neutralinos, keeping a correct relic density. This can be achieved, provided that χ_1^0 is mostly bino, or at least it possesses a substantial bino-component, and that there is an additional mechanism to

² It is not clear at the moment if the pure-wino case is consistent with DM indirect detection [12–14], due to the large uncertainties involved.

increase $\langle\sigma_{\text{ann}}v\rangle$. There are three of such mechanisms, which have been extensively studied in the literature:

- i)* Well-tempered neutralinos. If the parameters of the \mathcal{M}_{χ^0} matrix are finely chosen, χ_1^0 may be a well-tempered neutralino [15], i.e. an appropriate mixture of bino and Higgsino (or bino, Higgsino and wino), such that it annihilates in the right amount at early times. Since the $\propto M_Z$ off-diagonal entries in \mathcal{M}_{χ^0} are typically much smaller than M_1, M_2 and μ , a significant mixing requires some of the latter parameters to be near-degenerate.
- ii)* Funnels. If χ_1^0 is close to a bino, it can annihilate resonantly via Z -funnel, Higgs-funnel or A -funnel, provided its mass is nearly half of the mass of the funnel-particle.
- iii)* Co-annihilation. The effective $\langle\sigma_{\text{ann}}v\rangle$ increases if χ_1^0 can co-annihilate with other fast-annihilating particle (e.g. a stop, a stau or a gluino). This requires their masses to be nearly-degenerate.

In all the above cases, one can foresee the need of cancellations or delicate balances, and thereby fine-tuning.

The aim of this paper is to analyze all these possibilities in detail, evaluating the associated fine-tuning. In some cases, this requires to re-visit the concept of fine-tuning itself, because the extrapolation of the “standard criterion”, eq. (1.1), to the relic density is not always appropriate. The paper is organized as follows. In section 2, we will review the different measurements of the fine-tuning. Sections 3, 4, 5 and 6 are devoted to the different scenarios for DM within the MSSM. In section 7, we make a connection between the fine-tuning in DM and the electroweak fine-tuning. Section 8 is devoted to accommodating Higgsino DM in the MSSM and finally our conclusions are presented in section 9.

2 The measure of the fine-tuning

In the few places of the literature where the fine-tuning associated with the DM relic density has been considered, the criterion to quantify it has always been the standard one, i.e. a direct extrapolation of the EW fine-tuning criterion (1.1) replacing v^2 by Ω_{DM} ,

$$\Delta_i^{(\text{DM})} = \frac{d \log \Omega_{\text{DM}}}{d \log \theta_i} ; \quad \Delta^{(\text{DM})} \equiv \max \left\{ \Delta_i^{(\text{DM})} \right\}. \quad (2.1)$$

However, behind this “standard measure” there are implicit assumptions (seldom stated). If these assumptions do not hold, then the standard criterion may be misleading. In the next subsection, we compile those assumptions, and later we will show instances where those conditions are not fulfilled and therefore the standard criterion is not applicable. As we will see in the following sections, these instances are actually realized in several cases of DM production, which requires to improve the criterion to quantify the fine-tuning.

2.1 Assumptions behind the standard fine-tuning criterion

Let us now analyze the statistical meaning of the standard fine-tuning criterion, eq. (1.1). For the sake of simplicity, we will consider a single and representative θ -parameter, e.g. the one producing the maximum Δ (usually θ is a soft mass or the μ -parameter),

$$\Delta_\theta = \frac{\partial \log v^2}{\partial \log \theta}. \quad (2.2)$$




Figure 1. Schematic representation of the statistical interpretation of the standard fine-tuning criterion as the (inverse of the) p -value, $\Delta^{-1} = \delta\theta/\theta_0$.

As it is known, the issue of the value of v^2 is that it receives contributions of the size of the soft squared-masses, which are typically $\mathcal{O}(100)$ times larger; thus a somewhat artificial cancellation among these contributions is required. Since for non-tuned values of the soft terms (represented by θ), v^2 tends to be too large, one can estimate the small range of θ for which v^2 is abnormally small, say $v^2 \lesssim (v^{\text{exp}})^2$. Expanding $v^2(\theta)$ at first order around the value θ^0 , which gives $(v^{\text{exp}})^2$, $v^2(\theta_0 + \delta\theta) \simeq v^2(\theta_0) + (\partial v^2(\theta)/\partial\theta)_{\theta_0} \delta\theta$, we find that only for a small neighbourhood $\delta\theta \simeq \theta^0/\Delta_\theta$ around this point, v^2 is equal or smaller than the experimental value (see Fig. 1). Therefore, if one assumes that θ could reasonably have taken any value of the order of magnitude of θ^0 , then only for a small fraction $\sim |\delta\theta/\theta^0| \simeq \Delta_\theta^{-1}$ of the θ values one gets $v^2 \lesssim (v^{\text{exp}})^2$; this is the rough probabilistic meaning of Δ_θ [16, 17]. Consequently, Δ can be interpreted as the inverse of the p -value to get v^2 equal to the observed value or even smaller,

$$p\text{-value} \simeq \left| \frac{\delta\theta}{\theta_0} \right| \equiv \Delta^{-1}. \quad (2.3)$$

Then, we can summarize the implicit assumptions behind the standard fine-tuning criterion, eq. (1.1):

1. The possible values of a θ -parameter are distributed, with approximately flat probability, in the $\sim [0, \theta^0]$ range (flat prior in the Bayesian language). Note that, in fact, this represents two assumptions.
2. The expansion of $v^2(\theta)$ at first order captures its behaviour in the neighbourhood of interest.

If any of these assumptions is not fulfilled, then the standard criterion has to be re-visited. Before showing some typical examples where this can happen, let us add some comments on the above conditions.

The assumed range for θ does not need to be $[0, \theta_0]$, any range of the same length, e.g. $[\theta_0/2, 3\theta_0/2]$, works equally well. The idea is that the range for θ should be of the same order than its actual value, θ_0 , so that the latter is a typical value. It could be argued that in the upper half of the previous alternative range, i.e. $[\theta_0, 3\theta_0/2]$ it happens that $v^2 \leq (v^{\text{exp}})^2$, simply because $v^2 = 0$ for most of it. Then the p -value would be $\simeq 1/2$. Nevertheless, the region where v^2 is strictly vanishing should not be counted since it does not represent any extreme case but simply the case where the Higgs mass-squared parameter is positive. An equivalent way to take this fact

into account is to directly define the p -value for the mass parameter itself, m^2 , instead of v^2 (both are related by $v^2 = -m^2/\lambda$). Then, one evaluates the probability of having $|m^2| \leq |m^{\text{exp}}|^2$, giving a similar result as eq. (2.3).

The previous discussion illustrates the fact that there is always an $\sim \mathcal{O}(1)$ factor of arbitrariness for the fine-tuning measure. E.g. choosing the range of θ two times longer than the previous one increases Δ_θ by a factor 2.

It is also worth mentioning that the standard fine-tuning criterion is also valid for alternative choices of the prior and range of θ . E.g., if one assumes that θ has a logarithmic prior, i.e. its a-priori probability distribution is flat in the logarithm, $\mathcal{P}(\theta) \propto 1/\theta$, then a similar argument leads to the same eq. (2.3), provided that the range of θ satisfies $\log |\theta^{\text{max}}/\theta^{\text{min}}| = 1$.

Let us finally mention that the previous discussion about the statistical meaning of the fine-tuning can be expressed in Bayesian terms, following a Bayesian analysis of the probability distribution in the parameter space, see refs. [18, 19].

2.2 Examples

The EW fine-tuning stemming from the artificial cancellation between different contributions in order to get v^2 small enough, does reasonably fulfil conditions 1 and 2 of the previous subsection. In other words, in the MSSM the dependence of v^2 on the relevant soft terms and the μ -parameter goes as in Fig. 1 or behaves in a similar manner. So the standard criterion to quantify the EW fine-tuning is sound.

Now, let us suppose that the fine-tuned quantity, say F , has a different dependence on θ . Figs. 2 and 3 show two instances in which this happens in distinct ways. In Fig. 2, the hypothetical F -quantity acquires its experimentally observed value $F^{(\text{obs})}$ for some value θ_0 . However, there is no value of θ for which F vanishes. Hence, the region of $\delta\theta$ for which $F \leq F^{(\text{obs})}$ cannot be approximated by $\delta\theta \simeq \theta_0/\Delta_\theta$. The actual $\delta\theta$ region is narrower and thus the actual p -value is smaller and the fine-tuning is more severe. This example also illustrates another potential departure from the conditions 1 and 2 stated in the previous subsection. Obviously, if the value of θ_0 that reproduces $F^{(\text{obs})}$ lies *very* close to the minimum of the $F(\theta)$ function, then the fine-tuning is enormous, since essentially the $\delta\theta$ region for which $F \leq F^{(\text{obs})}$ shrinks to a point. Nevertheless, a blind application of the standard criterion would lead to $\Delta \rightarrow \infty$. Evidently, the problem is that in this case θ_0 would be a stationary point and thereby it would be no longer justified to truncate the expansion at first order (condition 2 in the previous subsection). An important lesson is that sensitivity is not always equivalent to fine-tuning, and sometimes the measure of sensitivity, which is what the standard criterion provides, does not reflect the actual degree of fine-tuning. As we will see, when the relic density gets the observed value thanks to the annihilation of neutralinos through Z , Higgs or A funnels, Ω_{DM} has a dependence on the MSSM parameters similar to that of Fig. 2.

Fig. 3 shows another example in which the assumptions for the applicability of the standard fine-tuning criterion do not hold. In this case, the truncation of $F(\theta)$ at first order is not good enough to evaluate the region $\delta\theta_0$ for which $F \leq F^{(\text{obs})}$. Here, the linear approximation leads to an underestimation of $\delta\theta_0$, so that the actual p -value is larger and the fine-tuning is less severe than the one obtained from the standard criterion. As we will see, the example of Fig. 3 describes schematically the dependence of Ω_{DM} on the MSSM parameters when the DM is wiped out through co-annihilations.

3 Well tempered bino-Higgsino

Consider first the well-tempered Higgsino/bino, i.e. the case in which the lightest neutralino is a combination of bino and Higgsino. Obviously, this scenario includes the pure-Higgsino case as




Figure 2. A hypothetical case where the standard criterion (see Fig. 1) underestimates the severity of the fine-tuning.




Figure 3. A hypothetical case where the standard criterion (see Fig. 1) overestimates the severity of the fine-tuning.

a particular and important limit (recall that, in contrast, the pure-bino limit is not viable unless additional mechanisms for DM annihilation are present). As mentioned in the introduction, the appeal of this setup is that it enables cases where the LSP is lighter than in the pure-Higgsino case, since the annihilation of LSPs becomes reduced thanks to the bino component. On the other hand, the possibility to find DM in (spin-independent) direct detection experiments through the neutralino elastic scattering off quarks mediated by a Higgs boson, is also higher, due to the Higgsino-bino-Higgs coupling. Indeed, present bounds on direct detection are able to exclude a large portion of the bino-Higgsino parameter space [6, 20]. However, it still remains as an interesting scenario, with relevant implications for the LHC and DM direct detection searches. It is also an illustrative example of the subtleties involved in the calculation of the DM fine-tuning.

From the four parameters that define the neutralino mass matrix, eq. (1.2), the most relevant ones here are M_1 and μ . M_2 plays a negligible role, unless it happens to be quite degenerate with M_1 and μ , in which case the neutral and charged wino would contribute to DM co-annihilation processes. This would correspond to the bino/wino/Higgsino scenario, to be analyzed in the next subsection. Consequently, for the bino-Higgsino analysis, M_2 can be made large enough for winos

to be ignored. In addition, as stated in the introduction, one needs $\tan\beta$ at least moderately large, say $\tan\beta \gtrsim 6$, in order to maximize the tree-level Higgs mass ($m_h^{\text{tree}} \leq M_Z$). In this way, we avoid the necessity of large radiative corrections to increase m_h up to its experimental value, which would require enormous stop masses and thereby an extremely large EW fine-tuning. Since the aim of this work is to explore as less fine-tuned as possible supersymmetric DM, we will ignore the small $\tan\beta$ regime. On the other hand, in the large- $\tan\beta$ regime the precise value of $\tan\beta$ is not very important, because it hardly affects the numerical values of the \mathcal{M}_{χ^0} entries. In conclusion, concerning the potential fine-tuning to arrange the correct DM relic density, $\tan\beta$ can be safely ignored.




Figure 4. Region of the $\mu - M_1$ plane that leads to the observed DM relic density, $\Omega_{\chi_1^0} h^2 = 0.119 \pm 0.012$, in a well-tempered bino-Higgsino scenario (blue bands). The grey band is excluded by LEP limits on charginos.

Fig. 4 shows in blue the region in the $\mu - M_1$ plane where $\Omega_{\chi_1^0} h^2 = 0.119 \pm 0.012$. It is located close to the $|\mu| = M_1$ lines, something required in order to get a non-trivial bino-Higgsino mixture. The calculation has been performed using `SOFTSUSY-3.6.2` [21] to compute the mass spectrum, `micrOMEGAs-4.1.8` [22, 23] for the relic density and direct detection cross section, and `MultiNest-3.9` [24–26] to efficiently explore the parameter space. The current LUX exclusion line [27] and the preliminary LUX 2016 limit [28] for the two signs of μ are presented in Fig. 5, showing the impressive power of present and future experiments of DM direct detection to exclude large regions of the parameter space. In fact, the (non-visible) XENON 1T and LZ projected sensitivities lie below the horizontal axes, so that they will potentially probe the whole scenario [29–31].

Let us now consider the DM fine-tuning issue. From Fig. 4 it is clear that a certain fine-tuning is required for the viability of the model, since the (low-energy) values of $|\mu|$ and M_1 must be quite degenerate. In the absence of a theoretical argument to justify such coincidence, this clearly represents a fine-tuning.

Before attempting to quantify it, let us mention an interesting and fortunate fact. The degree of naturalness of a physical scenario must be evaluated by examining the behaviour of the fine-tuned quantities with respect to the independent parameters of the theory, see e.g. the standard measure of eq. (1.1). Here, “independent” means that there is no known theoretical connection between them (or no connection based on some specific model is assumed). For the present case, the relevant independent parameters are the initial (high-energy) values of the soft parameters and μ . E.g. $\tan\beta$ is a derived parameter, which depends on the initial ones in a complicated way. Nevertheless, as mentioned above, the dependence of \mathcal{M}_{χ^0} on $\tan\beta$ is very weak, so we can




Figure 5. Spin-independent neutralino-proton cross section for the bino-Higgsino scenario for $\mu > 0$ (left) and $\mu < 0$ (right) and different values of $\tan \beta$. The current exclusion line and the preliminary 2016 limit from LUX (assuming that the neutralino is entirely made of bino-Higgsinos) are shown as solid and dashed lines, respectively. The XENON 1T and LZ projected sensitivities lie below the horizontal axes.

ignore its impact on the fine-tuning. Now, the fortunate fact is that the remaining three relevant (low-energy) parameters, involved in \mathcal{M}_{χ^0} , namely M_1, M_2 and μ , are essentially in one-to-one multiplicative correspondence with the three initial (high-energy) parameters,

$$\begin{aligned} M_i|_{LE} &= c_{M_i} M_i|_{HE}, \quad i = 1, 2, \\ \mu|_{LE} &= c_\mu \mu|_{HE}, \end{aligned} \quad (3.1)$$

where the HE (LE) subscript denotes high- (low-) energy, and the values of the c -coefficients depend on the value of the HE scale (see ref. [17] for a recent computation). However, for fine-tuning purposes the particular values of the c 's, and thus the choice of the HE scale, are irrelevant. E.g. for the standard fine-tuning measure, eq. (2.1), the logarithmic derivatives are the same evaluated with respect to the HE or the LE parameters. This fact simplifies life considerably and allows to work just with the low-energy parameters, producing results on $\Delta^{(DM)}$ which are pretty general, in particular $\Delta^{(DM)}$ is essentially independent of the HE scale and the values of the remaining MSSM parameters, which is remarkable. Incidentally, this is not the case for the EW fine-tuning, where a specific analysis must be performed for each model.

Let us now compute the DM fine-tuning. Before relying on the standard measure, eq. (2.1), it is convenient to test if the conditions 1 and 2 listed in subsection 2.1 are fulfilled. In other words, we should check the dependence of $\Omega_{\chi_1^0}$ on the μ and M_1 parameters (the only relevant ones for this scenario). Since the tuning is precisely between these two parameters, it is enough to consider one of them, say M_1 .³ Fig. 6 shows such dependence for a fixed value of μ .

As expected, only for a small interval of M_1 , $\Omega_{\chi_1^0}$ is consistent with the observed value. Nevertheless, concerning the fine-tuning, the important issue is that, *typically*, $\Omega_{\chi_1^0}$ is much larger or much smaller than $\Omega_{DM}^{(obs)}$. It requires a tuning between M_1 and μ for $\Omega_{\chi_1^0}$ to be in the vicinity of the observed value. Now, if we consider that the range of M_1 is $[0, M_1^{(0)}]$, where $M_1^{(0)}$ is the value

³This has the advantage of avoiding interference with the EW fine-tuning, for which μ is a very relevant parameter, unlike M_1 .




Figure 6. $\Omega_{\chi_1^0} h^2$ vs M_1 in the well-tempered bino-Higgsino scenario for fixed values of μ and $\tan \beta$.

that reproduces $\Omega_{\text{DM}}^{(\text{obs})}$, then the standard measure of eq. (2.1) and its interpretation in terms of p -value, i.e. the probability of getting $\Omega_{\chi_1^0} \leq \Omega_{\text{DM}}^{(\text{obs})}$, is justified. However, changing the limits of the range to e.g. $[M_1^{(0)}/2, 3M_1^{(0)}/2]$ jeopardizes the p -value interpretation, because there is a large interval of M_1 for which $\Omega_{\chi_1^0} \leq \Omega_{\text{DM}}^{(\text{obs})}$. A way out to this difficulty is to change the definition of the fine-tuned quantity. Instead of $\Omega_{\chi_1^0}$, we can use the mixing angle, θ , between the bino and the Higgsino. More precisely, upon diagonalization of \mathcal{M}_{χ^0} , given by eq. (1.2), one gets

$$|\tan 2\theta| \simeq \frac{\sqrt{2} s_W M_Z}{|\mu - M_1|} . \quad (3.2)$$

It is worth noting that θ is a physical quantity, in direct correspondence with $\Omega_{\chi_1^0}$, which could have been experimentally measured before $\Omega_{\text{DM}}^{(\text{obs})}$. If $\tan 2\theta$ is large, this clearly denotes a fine-tuning between M_1 and μ in eq. (3.2). In terms of $\tan 2\theta$ the p -value interpretation of the fine-tuning is much more transparent and robust than before: it is the probability of getting $|\tan 2\theta| \geq |\tan 2\theta^{(\text{obs})}|$. Assuming, as usual, a flat prior for M_1 in the region of interest, such p -value is simply

$$p\text{-value} = \frac{2|\mu - M_1|}{|M_1|} , \quad (3.3)$$

independent of the position of the M_1 -range limits.

Fig. 7 shows the fine-tuning calculated with the standard criterion eq. (2.1) and the one estimated by the inverse of the p -value, eq. (3.3).⁴ Needless to say, a $(p\text{-value})^{-1} = \mathcal{O}(1)$ is completely normal for a non-fine-tuned quantity, so fine-tunings below 5 or even 10 are not significant.

⁴Let us note the funny fact that if one had applied the standard fine-tuning criterion to the physical quantity $\tan 2\theta$ instead of $\Omega_{\chi_1^0}$, i.e. $\Delta = d \log \tan 2\theta / d \log M_1$, the result would have become essentially




Figure 7. Fine-tuning in the well-tempered bino-Higgsino scenario for $\mu > 0$, calculated using the “standard criterion” [see eq. (2.1)] (cyan band) and using the p -value criterion [i.e., the inverse of the p -value evaluated as in eq. (3.3)] (blue band). The width of the bands corresponds to the uncertainty in the relic density, $\Omega_{\chi_1^0} h^2 = 0.119 \pm 0.012$.

Qualitatively both criteria give similar results. In particular, the region around $M_1 = 500 - 600$ GeV is the most fine-tuned one, since it is the one that requires $\mu = M_1$ with more precision (it corresponds to maximal bino-Higgsino mixing angle). This can also be seen at naked eye in Fig. 4, by examining the width of the $\Omega_{\chi_1^0} h^2 = 0.119 \pm 0.012$ (blue) band, which narrows in that region. Quantitatively, the fine-tuning estimated by the p -value criterion is in general more severe and, in our opinion, more reliable for the above-discussed reasons. Interestingly, for the $M_1 \gtrsim 950$ GeV region, which is the one allowed by LUX, see Fig. 5, the tuning is rather small, even non-significant. This includes, of course, the $M_1 > \mu \simeq 1$ TeV region, for which the lightest neutralino is essentially a Higgsino. Actually, in this limit the precise value of M_1 is irrelevant, and the dependence of $\Omega_{\chi_1^0}$ on μ , namely $\Omega_{\chi_1^0} \propto \mu^2$, does not entail any fine-tuning, as expected, see the discussion in section 1.

4 Well-tempered bino-wino(-Higgsino)

By inspection of the \mathcal{M}_{χ^0} mass matrix, eq. (1.2), it is clear that a substantial bino-wino mixing requires a not too-large μ . Thus, assuming again moderate or large $\tan \beta$, this scenario has three relevant parameters, M_1 , M_2 and μ . Furthermore, the Higgsino gets also mixed, so that the scenario really becomes well-tempered bino-wino-Higgsino.

However, there is a special and physically relevant limit, where things become simpler. Namely, for large enough μ , the mixing between bino and wino (and Higgsino) is small. In that regime, pro-

equivalent to the inverse of the p -value, eq. (3.3). This shows that the standard criterion is not always robust under changes in the definition of the fine-tuned quantity. However, starting directly with the p -value criterion is much more trustworthy.

vided M_1 and M_2 are nearly-degenerate, the neutralino annihilation is dominated by co-annihilation with winos (more precisely, by wino-annihilation provided these are in thermal equilibrium with the lightest neutralino) [15] and is almost independent of the value of μ . All this is illustrated in Fig. 8, which shows the region in the $M_1 - M_2$ plane where $\Omega_{\chi_1^0} h^2 = 0.119 \pm 0.012$ for three different values of μ and the two signs of M_2 . For $|\mu| \gg |M_1|, |M_2|$ the solution is close to the straight band $|M_1| \simeq |M_2|$ and is quite independent of μ (the larger μ , the more independent the solution). This scenario can still be called well-tempered bino-wino, even though χ_1^0 is mostly bino.

In this figure, the $|\mu| \simeq |M_1| \ll |M_2|$ regions (nearly vertical segments of the coloured bands) are also visible. They correspond to the bino-Higgsino solution (analyzed in the previous section), and are quite independent of the value of $|M_2|$. Likewise, the $|\mu| \sim |M_1| \sim |M_2|$ regions (short, curved parts of the bands) correspond to the bino-wino-Higgsino case, to be discussed later.




Figure 8. Region of the $M_1 - M_2$ plane that leads to the observed DM relic density, $\Omega_{\chi_1^0} h^2 = 0.119 \pm 0.012$, for three different values of μ in a well-tempered neutralino scenario.

For the fine-tuning discussion of the bino-wino scenario, it is useful to consider some analytical approximations. Note that, since this is a co-annihilation scenario, the averaged annihilation cross-section $\langle \sigma_{\text{ann}} v \rangle$ in eq. (1.4) must be replaced by [15]

$$\langle \sigma_{\text{eff}} v \rangle = \frac{\sum_{i,j=1}^N w_i w_j \sigma_{ij} x^{-n}}{(\sum_{i=1}^N w_i^2)^2}, \quad w_i = \left(\frac{m_i}{m_1} \right)^{3/2} e^{-x \left(\frac{m_i}{m_1} - 1 \right)}, \quad (4.1)$$

where N is the number of co-annihilating species (in this case the bino and the three winos), m_1 is the lowest mass (in this case $\sim M_1$) and the $ij \rightarrow \text{SM SM}$ annihilation-cross-sections are parametrized as the dominant term in the velocity- (or equivalently x -) expansion

$$\langle \sigma_{ij} v \rangle \simeq \sigma_{ij} x^{-n}. \quad (4.2)$$

Under these circumstances the neutralino relic abundance is mostly determined by the \tilde{W} annihilation processes, whose cross sections go as $\sim g^4/M_2$. Plugging numerical factors one arrives

to a good approximate expression for the relic density [15],

$$\Omega_{\chi_1^0} h^2 \simeq 0.13 \left(\frac{M_2}{2.5 \text{ TeV}} \right)^2 \frac{1}{R_{\tilde{W}}}, \quad (4.3)$$

where

$$R_{\tilde{W}} = \int_0^1 dy \left[1 + \frac{1}{3} \left(\frac{M_1}{M_2} \right)^{3/2} e^{\frac{x_f}{y} \left(\frac{M_2}{M_1} - 1 \right)} \right]^{-2} \simeq \left(\frac{3}{4} \right)^2 e^{-\xi_{\tilde{W}} x_f \left(\frac{M_2}{M_1} - 1 \right)}, \quad (4.4)$$

with $\xi_{\tilde{W}} \simeq 1.7$. Recalling that $x_f \sim 20$, the previous equations (4.3, 4.4) show a strong sensitivity of $\Omega_{\chi_1^0}$ to M_1 . This is illustrated in Fig. 9, which shows $\Omega_{\chi_1^0} h^2$ vs. M_1 , using the complete numerical evaluation performed with micrOMEGAS, for fixed values of M_2 and $\tan \beta$. The value of μ is quite irrelevant provided is large enough ($\mu = 1.5$ TeV in the figure).




Figure 9. $\Omega_{\chi_1^0} h^2$ vs M_1 in the well-tempered bino-wino scenario for fixed $\mu = 1.5$ TeV, $M_2 = 500$ GeV and $\tan \beta = 10$. Left (right) panel shows the relic abundance in logarithmic (linear) units.

From eqs. (4.3, 4.4) and Fig. 9, one could foresee that the standard criterion will point to a severe fine-tuning. On the other hand, it should be noticed that the application of the standard recipe eq. (2.1) to eqs.(4.3, 4.4) leads to a value of the fine-tuning, Δ , that is essentially independent of the mass difference $\Delta m = ||M_2| - |M_1||$, as $\Omega_{\chi_1^0}$ is dominated by the Boltzmann (exponential) factor in a substantial range of M_1 -values. This is counter-intuitive, since logically the fine-tuning should be more severe when Δm is required to be smaller. Fig. 9 (right panel), which shows the dependence of $\Omega_{\chi_1^0} h^2$ on M_1 in a linear scale, clarifies the connection of the standard fine-tuning measure to the p -value in this case. Evidently, the exponential shape leads to an over-estimation of the fine-tuning when this is calculated with the standard criterion, compare Fig. 9 (right panel) to Fig. 3. Again, we find that the simple “ p -value-like measure”, $\Delta m/M_1$, offers a more sensible and robust description of the fine-tuning, as it happened in the bino-Higgsino case analyzed in section 3.

Fig. 10 shows the performance of both criteria. For each value of M_1 , the corresponding M_2 is chosen so that the observed relic density (1.3) is fulfilled (recall that the value of μ is large and fairly irrelevant). Both M_1 and M_2 are defined at the $Q = M_1$ scale, and their values are close to the physical masses, $m_{\chi_1^0}$ and $m_{\chi_2^0}, m_{\chi_1^\pm}$, respectively. As discussed above, the standard criterion leads to an almost flat fine-tuning, independently of M_1 and Δm . The p -value criterion, however, varies considerably with M_1 , showing a rather mild fine-tuning when the neutralino is light. The reason is that the heavier the wino, the less efficient its annihilation. Hence, in order to reproduce

the relic density, the Boltzmann penalty in the co-annihilation process must be lessened, which requires a smaller $\Delta m/M_1$, and thus a higher fine-tuning.




Figure 10. Fine-tuning in the well-tempered bino-wino scenario for $\mu = 1.5 \text{ TeV}$, $\tan \beta = 10$, calculated using the “standard criterion” [see eq. (2.1)] (cyan band) and using the p -value criterion, i.e. $(\Delta m/M_1)^{-1}$ (blue band). The width of the bands corresponds to the uncertainty in the relic density, $\Omega_{\chi_1^0} h^2 = 0.119 \pm 0.012$.

Interestingly, the fine-tuning (evaluated with the p -value criterion) is milder when $|\mu|$ approaches the value of $|M_1|$ or $|M_2|$. This is due to the fact that as $|\mu|$ decreases the mixing between bino and wino increases. Then, the neutralino annihilation does not only occur through wino co-annihilation, as explained above, but also through direct $\chi_1^0 \chi_1^0 \rightarrow \text{SM SM}$, $\chi_1^0 \chi_1^\pm \rightarrow \text{SM SM}$ processes, thanks to the non-negligible wino component of the neutralino. Since the resulting annihilation is now more efficient, $|M_1|$ does not need to be that close to $|M_2|$. Consequently, the p -value is larger (and the fine-tuning less severe). This effect is illustrated in Fig. 11, which displays the upper part of Fig. 8 (positive M_1/M_2 plane), but explicitly showing the percentage of $\chi_1^0 \chi_1^0 \rightarrow \text{SM SM}$ annihilation.

The situation depicted above connects with the bino-wino-Higgsino case, which occurs when the three relevant parameters, μ , M_1 and M_2 , have similar absolute values (curved segments of the bands in Fig. 11). Intuitively, this case requires a more severe fine-tuning, as it requires a “conspiracy” between three (a priori) independent parameters. It is therefore disfavoured from the point of view of naturalness, which is the main concern of this paper. One can try to estimate the related p -value. Assuming that μ is a given value, the separate p -values associated with the tuning of M_1 and M_2 are of order

$$\sim \left| \frac{|\mu| - |M_1|}{M_1} \right|, \quad \left| \frac{|\mu| - |M_2|}{M_2} \right|, \quad (4.5)$$

respectively. They should be combined multiplicatively. It is easy to check from Fig. 11 that this




Figure 11. As Fig. 8, zoomed on the positive region of the $M_1 - M_2$ plane. The color code denotes the percentage of $\chi_1^0 \chi_1^0 \rightarrow \text{SM SM}$ annihilation (the remaining DM annihilation proceeds mainly via co-annihilation with winos).

leads to a fine-tuning which is typically an $\mathcal{O}(1 - 10)$ factor more severe than the tuning in the regions related to well tempered bino-Higgsino or bino-wino, as expected.

5 Funnel

When the LSP can resonantly annihilate in the s -channel through an intermediate boson, say F , and $m_{\chi_1^0} \simeq m_F/2$, the annihilation cross-section increases enormously. In this way, scenarios of almost pure bino, which normally lead to excessive relic density, can be rescued. The funnel particle, F , can be the Z -boson, the ordinary Higgs boson, h , and the pseudoscalar, A . Note that for the first two cases χ_1^0 must be rather light, which implies, as a matter of fact, that it should be nearly pure bino; otherwise, either M_2 or μ would be necessarily close to $m_Z/2$ or $m_h/2$, thus leading to charginos below the LEP limit, $M_{\chi^\pm} \gtrsim 100$ GeV. For the A -funnel case, it is desirable that χ_1^0 be mostly bino as well, otherwise its mass should be very large. Notice here that, without the help of any funnel, the annihilation cross section of pure Higgsinos or pure winos is already quite efficient, which requires them to be rather heavy ($\simeq 1$ TeV and $\simeq 3$ TeV respectively) in order to reproduce the correct relic density. If, in addition, there is a channel of resonant annihilation (funnel), then their masses should be even larger, which would imply a heavier supersymmetric spectrum, and thus a more severe EW fine-tuning.

Let us first consider the A -funnel, i.e. the resonant annihilation through the A pseudoscalar

$$\chi_1^0 \chi_1^0 \rightarrow A \rightarrow \text{SM SM}, \quad (5.1)$$

where $\text{SM SM} = b\bar{b}, gg$, etc. Note that, for this process to take place, χ_1^0 must have a non-vanishing component of Higgsino, so that the $\chi_1^0 - \chi_1^0 - A$ vertex is in fact $\tilde{B} - \tilde{H}^0 - A$. Consequently, the larger the Higgsino component of χ_1^0 (and thereby the smaller μ), the more efficient the annihilation. Another point to keep in mind is that, even if M_1 is below the resonant value, i.e. $m_A/2 - M_1 > \Gamma_A$,

there can still be resonant annihilations, thanks to the thermal agitation in the early Universe, for some collisions the kinetic energy of the neutralinos can be large enough to reach $s \simeq m_A/2$. Of course, this amounts to a “Boltzmann penalty” for the averaged cross section. On the other hand, if $M_1 > m_A/2$, resonant annihilations are not possible. Then, the relic density, eq. (1.4), as a function of M_1 shows a characteristic asymmetric dependence on M_1 in the resonance-region.

All this is illustrated in Fig. 12 for $m_A = 800$ GeV. Note that “far” from the resonant point the




Figure 12. $\Omega_{\chi_1^0} h^2$ vs M_1 in the A -funnel scenario for $m_A = 800$ GeV, $\tan \beta = 20$ and three (positive) values of μ . The grey band denotes the observed relic abundance.

dependence of $\Omega_{\chi_1^0}$ on M_1 has an exponential shape due to the above-mentioned Boltzmann penalty. Thus we can expect that, similarly to what happened for the bino-wino co-annihilation scenario (section 4), the standard criterion for the fine-tuning is not suitable here and typically overestimates the real fine-tuning. Nevertheless, when M_1 approaches the resonant point, we expect the opposite (recall the discussion around Fig. 2). In the limiting case, in which the physical region, $\Omega_{\chi_1^0} \simeq \Omega_{\text{DM}}^{(\text{obs})}$, is close to the minimum of the curve (this occurs for $\mu \simeq 4000$ GeV in the example of Fig. 12), then the standard criterion indicates no fine-tuning at all; however this is obviously the most fine-tuned case!

In contrast, the p -value criterion is very transparent and easy to apply. Let us call $M_1^{(0)}$ the value of M_1 that, for given m_A and μ , leads to $\Omega_{\chi_1^0} = \Omega_{\text{DM}}^{(\text{obs})}$. Then, only in the narrow range $M_1 \in [M_1^{(0)}, M_1^{(0)} + \Delta m]$, with $\Delta m \simeq |m_A/2 - |M_1^{(0)}||$, the relic density will be equal or smaller than the observed value, as it is clear from Fig. 12. Therefore, the p -value is (once more) simply $\Delta m / |M_1^{(0)}|$.

Fig. 13 illustrates the previous discussion, showing the fine-tuning calculated with the standard criterion, eq. (2.1), and the one estimated as the inverse of the p -value, eq. (3.3), for $m_A = 800$ GeV. For each value of M_1 , the corresponding $\mu > 0$ is chosen so that the observed relic density (1.3) is fulfilled (recall that the value of μ determines the amount of Higgsino mixing).

It is also worth-mentioning that this scenario is quite safe with respect to the current DM direct detection bounds because the elastic scattering cross section does not benefit from any resonant




Figure 13. Fine-tuning in the A -funnel scenario for $m_A = 800$ GeV and $\tan \beta = 20$, calculated using the “standard criterion” [see eq. (2.1)] (cyan band) and using the p -value criterion, i.e. $(|m_A/2 - |M_1||/|M_1|)^{-1}$ (blue band). For each value of M_1 , the corresponding (positive) μ is chosen so that $\Omega_{\text{DM}}^{(\text{obs})}$ is reproduced.

enhancement. This is shown in Fig. 14 for three different values of m_A .

Let us finally mention that for $|M_1| \lesssim \frac{3}{4}(m_A/2)$ the enhancement due to the resonant annihilation is lost, but the relic density can still be reproduced if μ (and/or M_2) are close enough to M_1 for the scenario to become a well-tempered neutralino case. In that case, the fine-tuning is due to this well-tempered character and has been analyzed in sections 3 and 4.

Now, we turn to the h - and Z -funnels

$$\chi_1^0 \chi_1^0 \rightarrow h \rightarrow SM \ SM, \quad (5.2)$$

$$\chi_1^0 \chi_1^0 \rightarrow Z \rightarrow SM \ SM. \quad (5.3)$$

Similarly to the A -funnel case, these channels require χ_1^0 to have a non-vanishing Higgsino component, so that the $\chi_1^0 - \chi_1^0 - h$ ($\chi_1^0 - \chi_1^0 - Z$) vertex has the $\tilde{B} - \tilde{H}^0 - h$ ($\tilde{H}^0 - \tilde{H}^0 - Z$) structure⁵. Thus, again, the larger the Higgsino component of χ_1^0 (and thus the smaller μ), the more efficient the annihilation. This effect is stronger for the Z -funnel, as it involves the Higgsino component in the two incoming neutralinos. All this is illustrated in Fig. 15, which shows the dependence of $\Omega_{\chi_1^0}$ vs M_1 for three different values of μ .

Regarding the fine-tuning issue, as for the A -funnel case, we expect that typically, the standard criterion overestimates the fine-tuning, except when the physical region, $\Omega_{\chi_1^0} \simeq \Omega_{\text{DM}}^{(\text{obs})}$, is close to a stationary point. Looking at Fig. 15, there are now three stationary points, corresponding to the two minima at $m_{\chi_1^0} \simeq m_h/2, M_Z/2$ and to the maximum between both. For the very same reasons

⁵For the h -funnel $\sigma_{\text{ann}} \propto |N_{11}N_{14}|^2$, while for the Z -funnel $\sigma_{\text{ann}} \propto |N_{11}N_{13} - N_{11}N_{14}|^4$, with N being the neutralino mass matrix.




Figure 14. Spin-independent neutralino-proton cross section in the A -funnel scenario for $m_A = 800$ GeV. The solid and dashed black lines denote the current LUX upper limit and the preliminary LUX 2016 bound, respectively. The XENON 1T (dash-dotted line) and LZ (dotted line) projected sensitivities are also depicted.




Figure 15. $\Omega_{\chi_1^0} h^2$ vs M_1 in the region of h - and Z -funnels for $\tan \beta = 10$ and three (positive) values of μ . The grey band denotes the observed relic abundance.

as for the A -channel, we find that a more robust and reliable measure of the fine-tuning is provided by the p -value $\simeq \Delta m/|M_1|$, where Δm is the length of the M_1 -range where the relic density is equal or smaller than the observed value⁶.

Fig. 16 (left panel), which is similar to Fig. 13 but for the h - and Z -funnels, illustrates the previous discussion. Again, for each value of M_1 , the corresponding $\mu > 0$ is chosen so that the observed relic density (1.3) is reproduced. As argued above, the standard criterion overestimates (non-dramatically) the fine-tuning in most of the M_1 -range. However, around the three stationary points (in particular the two resonant points), it underestimates the fine-tuning dramatically. Indeed, apart from the resonant points, the fine-tuning (estimated with the p -value criterion) is quite mild ($\lesssim 10$). Note from Fig. 15 that for 150 GeV $\lesssim \mu \lesssim 450$ GeV both the h -funnel and the Z -funnel can successfully reproduce a relic density equal to the observed size, or even smaller, if M_1 is positive and has the appropriate value. Thus, the two p -values should be added, implying a less severe fine-tuning. The results are shown in Fig. 16 (right panel). Notice that in this way the peak associated with the Z -funnel region is blown-up. This occurs because, for a given value of (positive) μ , the possibility of Z -funnel annihilation is always accompanied by the possibility of h -funnel annihilation, but not the other way round, see Fig. 15. For $\mu < 0$ the results are similar, but in that case the blown-up peak corresponds to the h -funnel for analogous reasons. The bottom line is that, apart from the peaks very close to the resonant points, the h - and Z -funnels show very mild or non-significant fine-tuning.




Figure 16. Left panel: Fine-tuning in the region of h - and Z -funnels for $\tan \beta = 10$, calculated using the “standard criterion” [see eq. (2.1)] (cyan band) and using the p -value criterion, as explained in the text (blue band). Right panel: The same but adding up the p -values corresponding to the h - and Z -funnels, when the value of μ allows for both possibilities (see text).

6 Annihilation and Co-annihilation

Co-annihilation occurs when one or several particles with masses close to the LSP annihilate efficiently. In that case, the relic density is still given by eq. (1.4), but with the, effective, averaged annihilation cross-section, $\langle \sigma_{\text{eff}} v \rangle$, given by eq. (4.1). Once more, this mechanism is only useful if

⁶ Around the Higgs-resonance $\Delta m \simeq |m_h/2 - m_{\chi_1^0}|$, while around the Z -resonance $\Delta m \simeq 2|m_Z/2 - m_{\chi_1^0}|$, due to the larger width of the Z -boson. This can be appreciated in Fig. 15.

the LSP is essentially a bino, which is the instance where the LSP does not annihilate efficiently enough at early times.

Due to the Boltzmann factor in eq. (4.1), $\langle\sigma_{\text{eff}}v\rangle$, and thus $\Omega_{\chi_1^0}$, is exponentially sensitive to the mass gap, Δm , between the neutralino and its neighbouring particles. Hence, the most important dependence on M_1 goes, qualitatively, as

$$\Omega_{\chi_1^0} \sim e^{-\xi x_f \left(\frac{\Delta m}{M_1}\right)}, \quad (6.1)$$

where $x_f \simeq 20$ and ξ is typically $\mathcal{O}(1)$. Some particularly important possibilities for the co-annihilating particles are the gluino, the stop and the stau (beside Higgsinos and winos, analyzed in previous sections). The above exponential dependence makes the standard criterion of fine-tuning to be quite severe in all cases,

$$\Delta_{M_1} = \left| \frac{\partial \log \Omega_{\chi_1^0}}{\partial \log M_1} \right| \simeq \xi x_f \frac{\tilde{m}}{M_1} = \mathcal{O}(1) \times 20, \quad (6.2)$$

where \tilde{m} is the mass of the co-annihilating particle. The puzzling, and suspicious, fact is that this estimation of the tuning does *not* depend on the mass difference $\tilde{m} - M_1$. It is essentially constant independently of how precisely M_1 should be close to \tilde{m} . Certainly, this is due to the fact that the standard criterion measures sensitivity rather than fine-tuning, and these are not always equivalent.




Figure 17. $\Omega_{\chi_1^0} h^2$ vs M_1 in the gluino co-annihilation scenario for fixed $m_{\tilde{g}} = 1200$ GeV. The lightest neutralino is essentially bino. Left (right) panel shows the dependence in logarithmic (linear) units.

In order to illustrate these aspects, let us consider the case of gluino co-annihilation. Fig. 17 (left panel) shows $\Omega_{\chi_1^0}$ vs M_1 for $m_{\tilde{g}} = 1200$ GeV. The exponential dependence on M_1 in the co-annihilation region has been zoomed in linear scale in the right panel. Now, comparing this figure to Fig. 3, it is clear that the standard criterion leads to an overestimation of the fine-tuning, since the truncation of $\Omega_{\chi_1^0}(M_1)$ at first order around the physical point is not good enough to describe the whole region where $\Omega_{\chi_1^0} \leq \Omega_{\text{DM}}^{(\text{obs})}$. Once again a more sensible measure is given by the p -value,

$$p\text{-value} \simeq \frac{\Delta m}{|M_1|}. \quad (6.3)$$

Fig. 18 shows the fine-tuning calculated with the standard criterion eq. (2.1) and the one estimated as the inverse of the p -value eq. (6.3) for gluino co-annihilation. For each value of M_1 , the

corresponding $m_{\tilde{g}}$ is chosen so that the observed relic abundance (1.3) is reproduced. As expected, the standard criterion clearly overestimates the fine-tuning and is suspiciously independent of M_1 . On the contrary, the p -value criterion shows a less severe tuning, especially for $M_1 \lesssim 500$ GeV, where it becomes almost irrelevant. The increase in this fine-tuning with M_1 occurs because the heavier the gluino, the less efficient becomes its annihilation, and this must be compensated by a more precise gluino-bino degeneracy.




Figure 18. Fine-tuning in the gluino co-annihilation scenario, calculated using the “standard criterion” [see eq. (2.1)] (cyan band) and using the p -value criterion, i.e. $(\Delta m/M_1)^{-1}$ (blue band). The $m_{\tilde{g}}$ value is chosen so that the observed relic density, $\Omega_{\tilde{\chi}_1^0} h^2 = 0.119 \pm 0.012$, is always reproduced. The lightest neutralino is essentially bino.

Other co-annihilation cases, as the ones mentioned above, show a similar pattern.

7 Connection to the electroweak fine-tuning

The DM fine-tuning must be combined with the EW one, since both affect the same theoretical scenario, namely the MSSM. We recall here that the EW fine-tuning is reasonably well estimated by the “standard measure”, $\Delta^{(\text{EW})} = \max \{ \partial m_h^2 / \partial \theta_i \}$ ⁷, where θ_i are the independent parameters of the model. As discussed in section 2 [see eq. (2.3)], this measure can be interpreted as the p -value associated with the small size of the EW scale. In sections 3–6, we have evaluated the analogous p -value to reproduce $\Omega_{\text{DM}}^{(\text{obs})}$. The main (computational) difference with the EW fine-tuning is that in this case the “standard measure” of the fine-tuning is not a reliable estimation of the p -value, so the latter has to be evaluated in a more direct way, as we have done. Typically, the EW fine-tuning is $\gtrsim \mathcal{O}(100)$, i.e. it is $\mathcal{O}(10)$ times more severe than the DM one, though the latter can be extremely larger at special places, see e.g. Figs. 7, 13 and 16. On the other hand, due

⁷This expression is equivalent to the expression (2.2), once the radiative corrections to the Higgs effective potential are taken into account [17].

to their common statistical interpretation (as p -values), it is clear that both fine-tunings should be multiplicatively combined. A subtle aspect here is that the EW and DM fine-tunings arise from cancellations between the same set of parameters.

This issue was analyzed in ref. [32], appendix A⁸. The idea is that when one computes the fine-tuning in a quantity, say Ω_{DM} , one is free to vary the input θ_i 's only in a way that all the potential constraints (in this case the EW scale) are fulfilled. Denoting, for simplicity, $(\Delta_{\text{EW}})_i$, $(\Delta_{\text{DM}})_i$ the EW and DM fine-tunings (i.e. the inverse p -values) with respect to the θ_i parameters, and $G(\theta_i) = 0$ the EW condition, one should project $\vec{\Delta}_{\text{DM}}$ into the subspace orthogonal to the $G(\theta_i) = 0$ hypersurface in the $\{\log \theta_i\}$ space. In other words, one has to re-define the DM fine-tuning as

$$\vec{\Delta}_{\text{DM}} \rightarrow \vec{\Delta}_{\text{DM}} - \frac{1}{|\vec{\Delta}G|^2} (\vec{\Delta}_{\text{DM}} \cdot \vec{\Delta}G) \vec{\Delta}G, \quad (7.1)$$

where $\vec{\Delta}G \equiv \{\partial G / \partial \log \theta_i\} \propto \vec{\Delta}_{\text{EW}}$.

From this discussion, it is clear that only the parameters that contribute substantially to both the EW and DM fine-tunings are to play a relevant role in the previous projection. In this sense, the EW fine-tuning is dominated by the initial values of $m_{\tilde{t}}, M_3$ and μ parameters (see e.g. ref. [17]), while the DM fine-tuning is dominated by M_1 , and, depending on the annihilation mechanism, by μ , M_2 , m_A or $m_{\chi'}$, where χ' is a possible co-annihilating particle (gluino, stop, etc.). Consequently, only the DM fine-tuning associated with μ or $m_{\chi'}$ is subject to be lowered by the non-trivial “interference” with the EW one. The conclusion is that the DM fine-tuning with respect to M_1 (which is the one computed in previous subsections) is always representative of the total DM fine-tuning and does not need to be corrected by the projection onto the subspace satisfying the EW condition.

As discussed in section 3, the DM fine-tuning is quite independent of the details of the MSSM scenario (whether it is CMSSM, NUHM, etc., or the value of the high-energy scale, M_{HE}). In this respect, it is a very robust feature of the MSSM. This fortunate circumstance does not occur for the EW fine-tuning: $\Delta^{(\text{EW})}$ is much more model-dependent, since it depends on the initial values of $M_3, m_{\tilde{t}}$, etc., and on the correlations between them (e.g. whether or not there is a universal scalar mass). It also depends on M_{HE} . Concerning this point, we can presume that $M_{\text{HE}} = M_X$, because a lower value for M_{HE} would typically lead to a very light gravitino⁹ ($m_{3/2} \sim (M_{\text{HE}}/M_{\text{P}})m_0$), which would then play the role of the LSP, instead of the lightest neutralino, as assumed in this paper. In any case, it is clear that for every scenario for which the DM fine-tuning has been computed in sections 3–6, there is not a unique value of $\Delta^{(\text{EW})}$; the latter depends on the details of the high-energy theory.

Nevertheless, instead of $\Delta^{(\text{EW})}$ one can consider $\Delta_{\text{min}}^{(\text{EW})}$, i.e. the minimal EW fine-tuning. Normally $\Delta^{(\text{EW})}$ is dominated by the M_3 -contribution or by the μ -contribution¹⁰. So one could just set $m_{\tilde{g}}$ at its experimental lower bound, ~ 1.3 TeV [34], which amounts to $M_3 \simeq 0.59$ TeV. This gives $\Delta_{M_3}^{(\text{EW})} = \mathcal{O}(100)$ [17], independently of the DM scenario, provided it can accommodate such light gluino. However, this is not always the case, e.g., as mentioned in the introduction, for pure-wino DM (which does not entail DM fine-tuning), $M_{\chi_1^0} \sim m_{\tilde{W}} \simeq 3$ TeV. This necessarily implies a heavy gluino, $m_{\tilde{g}} > 3$ TeV and, in turn, a much larger EW fine-tuning, near $\mathcal{O}(1000)$ (notice here that, parametrically, $\Delta_{M_3}^{(\text{EW})} \propto M_3^2 \propto m_{\tilde{g}}^2$). If the gaugino masses are unified at high energy,

⁸The same prescriptions were independently found in later references [4, 6].

⁹The gravitino could be heavier than this naive expectation in theories with extra dimensions, where gravity is stronger, see e.g. ref. [33].

¹⁰We are not considering here an (unknown) hypothetical scenario where all the soft terms and μ are theoretically correlated in such fortunate way that their contributions to m_h^2 nearly cancel, so that there is no fine-tuning!

the EW tuning is even larger, since $m_{\tilde{g}} \simeq 2.8m_{\tilde{W}}$ [17]. Other DM scenarios that may demand a heavy gluino are co-annihilation (with a particle different from gluino), A -funnel or well-tempered bino-wino-Higgsino, whenever $M_{\chi_1^0} \gtrsim 1.3$ TeV. In contrast, if the co-annihilation is with a gluino, this can be substantially lighter than 1.3 TeV, because it would be quite degenerate with the LSP, and thus invisible at the LHC. Hence, the latter scenario would reduce $\Delta_{M_3}^{(\text{EW})}$! Actually, from Fig. 18, we see that co-annihilation with a light gluino could improve both the DM and the EW fine-tunings.

Regarding the μ -contribution, this is $\Delta_\mu^{(\text{EW})} \simeq (2\mu/m_h)^2$ (for further details see ref. [17]), so for large enough μ , $\Delta^{(\text{EW})}$ becomes dominated by $\Delta_\mu^{(\text{EW})}$. More precisely, this happens for $\mu \gtrsim (1/2.3)m_{\tilde{g}}$, in particular for $\mu \gtrsim 570$ GeV if $m_{\tilde{g}}$ is close to its 1.3 TeV lower bound [34]. Consequently, the above-mentioned DM scenarios that implied a large $\Delta_{M_3}^{(\text{EW})}$, imply also an even larger $\Delta_\mu^{(\text{EW})}$. It is also worth mentioning that for pure Higgsino DM (which does not amount to DM fine-tuning), $M_{\chi_1^0} \sim \mu \simeq 1$ TeV, implying $\Delta_\mu^{(\text{EW})} \gtrsim \mathcal{O}(200)$.

From the previous discussion, it is clear that $\Delta^{(\text{DM})}$ should be kept as small as possible, preferably compatible with a non-fine-tuned situation, otherwise the combined fine-tuning will be above several thousands. This can be achieved in an obvious way if the DM is pure Higgsino or pure wino. Nonetheless, as mentioned, in the latter case the EW fine-tuning raises to $\gtrsim \mathcal{O}(1000)$ (for pure Higgsino it also grows but in a much milder way). Other cases that essentially imply no (or very mild) DM fine-tuning are: well-tempered bino-Higgsino (if $M_{\chi_1^0}$ is not around 500 GeV); Higgs, Z and A funnels when $M_{\chi_1^0}$ is not too close to (half) the resonance mass; and co-annihilation scenarios when $M_{\chi_1^0}$ is rather light, i.e. $\lesssim 500$ GeV.

Finally, in all the cases one has to ensure that *i*) the value of $m_{H_u}^2$ at LE has the right size ($\sim \mu^2$) to enable the correct EW breaking, and *ii*) the physical Higgs mass, $m_h \simeq 125$ GeV, is reproduced. Both facts have to do with the values of $m_{\tilde{t}_L}^2$, $m_{\tilde{t}_R}^2$, A_t , and $m_{H_u}^2$ at HE. In general, it will be possible to arrange these parameters so that (in combination with $M_3 \simeq 1.3$ TeV) they implement *i*) and *ii*) without significantly affecting the value of $\Delta^{(\text{EW})}$. However, if there are theoretical correlations between the initial soft terms, the *i*) and *ii*) conditions may imply further constraints on the theory and thereby an increase in $\Delta^{(\text{EW})}$. Next, we illustrate this point by considering the case of pure Higgsino DM (one of the preferred ones from the above discussion) when the theory is some kind of constrained MSSM.

8 Accommodating Higgsino DM in the MSSM

If the LSP is close to a pure Higgsino (with mass $\simeq \mu \simeq 1$ TeV), the rest of the supersymmetric particles must be heavier, which imposes conditions on HE parameters. Assuming $M_{\text{HE}} = M_X$ in what follows, one gets [17]

$$\begin{aligned} m_{\tilde{g}} &\simeq 2.22M_3 > 1 \text{ TeV} , \\ m_{\tilde{W}} &\simeq 0.8M_2 > 1 \text{ TeV} , \\ m_{\tilde{B}} &\simeq 0.43M_1 > 1 \text{ TeV} , \end{aligned} \tag{8.1}$$

where M_i are the gaugino masses at the HE scale. If these are unified, $M_1 = M_2 = M_3 \equiv M_{1/2}$, then from the last equation

$$M_{1/2} \gtrsim 2.3 \text{ TeV} , \tag{8.2}$$

implying

$$m_{\tilde{g}} > 5.16 \text{ TeV} . \tag{8.3}$$

This large value of $M_{1/2}$ implies a huge EW fine-tuning, see eq. (8.7) below. Other supersymmetric masses are also forced to be very large, e.g. the average stop mass, $m_t^2 \equiv \frac{1}{2}(m_{t_1}^2 + m_{t_2}^2)$, reads [17]

$$\overline{m}_t^2 \simeq \frac{1}{2}(5.945M_3^2 + 0.679m_{t_L}^2 + 0.611m_{t_R}^2 + 0.182M_2^2 - 0.307m_{H_u}^2 \cdots) + m_t^2. \quad (8.4)$$

Therefore, assuming gaugino unification leads to $m_{\tilde{t}} \gtrsim 4$ TeV. Similarly, using the formulae from ref. [17], one gets $m_{\tilde{L}} \gtrsim 1.5$ TeV. The singlet sleptons are much less constrained, but they are forced anyhow to live above 1 TeV so that the Higgsino plays the LSP role. All the previous relations are much less restrictive if one gives up gaugino unification or considers lower values of M_{HE} , though the latter possibility is disfavoured if the LSP is not the gravitino.

One can consider now the EW minimization condition, which, at (moderately) large $\tan \beta$, reads

$$-\frac{m_h^2}{2} = \mu^2 + m_{H_u}^2, \quad (8.5)$$

with all quantities defined at LE, in particular (see table 3 in appendix A of [17])

$$\begin{aligned} m_{H_u}^2|_{LE} = & -1.6M_3^2 + 0.63m_{H_u}^2 - 0.37m_{t_L}^2 - 0.29m_{t_R}^2 \\ & + 0.28A_tM_3 + 0.2M_2^2 - 0.13M_2M_3 - 0.11A_t^2 + \cdots, \end{aligned} \quad (8.6)$$

where the variables in the r.h.s. are at HE. Notice from eq. (8.5) that for $|\mu| \simeq 1$ TeV, $m_{H_u}^2 \simeq -|\mu|^2 \simeq -1$ TeV², which is an additional constraint. In fact, in the popular constrained-MSSM (CMSSM), it seems impossible to satisfy this constraint with all the supersymmetric masses higher than 1 TeV. Note that for the CMSSM, the contributions from $m_{H_u}^2$, $m_{t_L}^2$ and $m_{t_R}^2$ almost cancel in eq. (8.6), which is the well-known focus-point behaviour. Then, it is almost impossible to compensate the huge negative contribution coming from M_3^2 (recall that in the CMSSM $M_3 = M_{1/2} \simeq 2.3$ TeV, due to gaugino unification). Using a very large A_t with the appropriate sign does not help since the negative contribution from A_t^2 would dominate. Incidentally, pure wino DM is also unattainable, because whenever there is gaugino unification, the bino is lighter than the wino, so the latter cannot be the LSP.

Therefore, one has to go beyond the CMSSM. The non-universal-Higgs-Mass model (NUHM) is like the CMSSM, but allowing the soft Higgs masses, $m_{H_u}^2, m_{H_d}^2$ to be different from the other scalar masses at HE (a usual choice is $m_{H_u}^2 = m_{H_d}^2$). Then, if $m_{H_u}^2$ is large enough at HE, one can achieve $m_{H_u}^2(LE) \simeq -1$ TeV² in eq. (8.6). This implies that the extra Higgs states are quite heavy (if the $m_{H_u}^2 = m_{H_d}^2$ condition is imposed). The whole spectrum seems beyond the LHC. In addition, the model presents a high EW fine-tuning:

$$\Delta_{M_{1/2}}^{(\text{EW})} = \left| \frac{d \log m_h^2}{d \log M_{1/2}} \right| \simeq \left| -4 \frac{M_{1/2}^2}{m_h^2} \left(-1.6 + 0.2 - 0.13 + 0.14 \frac{A_t}{M_{1/2}} + \cdots \right) \right| \simeq 2000. \quad (8.7)$$

Another possibility is to start with non-universal gaugino masses. This is a much more flexible scenario and, in principle, it does not seem difficult in this case to achieve the LSP condition for the Higgsino and the correct EW breaking with supersymmetric masses (in particular gluino masses) not far from their experimental lower bounds.

9 Conclusions

One of the most celebrated bonuses of supersymmetric theories is the presence of stable WIMPs, which are natural candidates for dark matter (DM). In the MSSM, such role is usually played by the lightest supersymmetric particle (LSP), which is typically the lightest neutralino. However,

when one goes into the details, it turns out that in most scenarios some kind of tuning is needed in order to get Ω_{DM} of the right magnitude. This fine-tuning is worrisome since it has to be combined with the ubiquitous electroweak (EW) fine-tuning problem, i.e. the delicate balance between soft terms required to reproduce the smallness of the EW scale.

Taking into account that the original motivation for low-energy SUSY was to solve the hierarchy problem, which is the EW fine-tuning problem of the SM, it is logical to demand SUSY scenarios to be as natural as possible. In this sense, there exists a vast literature examining the EW fine-tuning problem, but little concerning the DM one.

In this paper, we study this problem in an, as much as possible, exhaustive and rigorous way. We have considered the MSSM framework, assuming that the LSP is the lightest neutralino, χ_1^0 , and explored various possible scenarios. These include different masses and compositions of χ_1^0 , which are completely defined by the parameters involved in the neutralino mass matrix ($M_1, M_2, \mu, \tan \beta$), as well as different mechanisms for neutralino annihilation in the early Universe (well-tempered neutralinos, funnels and co-annihilation scenarios). We also present a discussion about the statistical meaning of the fine-tuning and how it should be computed for the DM relic abundance, and combined with the EW fine-tuning. It turns out that the “standard measurement” of fine-tuning, $\Delta = d \log \Omega_{\text{DM}} / d \log \theta$ is not appropriate in most of the cases, and one has to evaluate the p -value associated with the smallness of Ω_{DM} , which, actually, amounts normally to a simpler computation. A fortunate fact is that the relevant (low-energy) parameters, involved in the neutralino mass matrix are essentially in one-to-one multiplicative correspondence with the initial (high-energy) parameters. This allows to compute the fine-tuning directly on the low-energy parameters with full generality. In consequence, the DM fine-tuning is quite independent of the details of the MSSM scenario (whether it is CMSSM, NUHM, etc., or the value of the high-energy scale). In this sense, it is a very robust feature of the MSSM. In contrast, the EW fine-tuning is much more model-dependent.

Concerning the results, the fine-tuning related just to the DM relic abundance is negligible or very mild in a number of scenarios. More precisely, when χ_1^0 is essentially a pure Higgsino or a pure wino there is no fine-tuning associated with the DM relic density. Other cases that essentially imply no (or very mild) DM fine-tuning are: well-tempered bino-Higgsino (if $M_{\chi_1^0}$ is not around 500 GeV); Higgs, Z and A funnels when $M_{\chi_1^0}$ is not too close to (half) the resonance mass; and co-annihilation scenarios when $M_{\chi_1^0}$ is rather light, i.e. $\lesssim 500$ GeV.

Nevertheless, this is not the end of the story, as the DM fine-tuning must be combined with the EW one. Modulo some subtleties discussed in this paper, both fine-tunings should be essentially multiplicatively combined. Thus, one should demand $\Delta^{(\text{EW})}$ to be as mild as possible. Normally $\Delta^{(\text{EW})}$ is dominated by the M_3 -contribution or by the μ -contribution. So one could just set $m_{\tilde{g}}$ at its experimental lower bound, ~ 1.3 TeV, which leads to $\Delta^{(\text{EW})} = \mathcal{O}(100)$, independently of the DM scenario, provided it can accommodate such light gluino [17]. However, this is not always the case. E.g. for pure-wino DM (which does not entail DM fine-tuning), $M_{\chi_1^0} \sim m_{\tilde{W}} \simeq 3$ TeV. This necessarily implies a heavier gluino and, in turn, a much larger EW fine-tuning, near $\mathcal{O}(1000)$. By contrast, if the co-annihilation is with a gluino, the latter can be substantially lighter than 1.3 TeV, since it would be invisible at the LHC. Hence, the latter scenario would reduce $\Delta^{(\text{EW})}$!

As a final remark, naturalness is a reasonable guide to look for plausible supersymmetric scenarios. In this regard, a strong emphasis has been put on the EW fine-tuning, but the DM fine-tuning is also very important, as shown in this paper, especially when it is combined with the EW one. This feature should be taken into account when one explores “natural SUSY” scenarios and their possible signatures at the LHC and in DM detection experiments.

Acknowledgements

This work has been partially supported by the MICINN, Spain, under contract FPA2013-44773-P, Consolider-Ingenio CPAN CSD2007-00042, as well as MULTIDARK CSD2009-00064. We also thank the Spanish MINECO *Centro de excelencia Severo Ochoa Program* under grant SEV-2012-0249. The work of A.D. is partly supported by the National Science Foundation under grant PHY-1520966. S.R. is supported by the Campus of Excellence UAM+CSIC. The work of M.E.C. was supported by Fundação de Amparo à Pesquisa do Estado de São Paulo. R. RdA is supported by the Ramón y Cajal program of the Spanish MICINN and also thanks the support by the “SOM Sabor y origen de la Materia” (FPA2011-29678), the “Fenomenologia y Cosmologia de la Fisica mas alla del Modelo Estandar e Implicaciones Experimentales en la era del LHC” (FPA2010-17747) MEC projects, the Severo Ochoa MINECO project SEV-2014-0398, the Consolider-Ingenio 2010 programme under grant MULTIDARK CSD2009-00064 and by the Invisibles European ITN project FP7-PEOPLE-2011-ITN, PITN-GA-2011-289442-INVISIBLES.

References

- [1] J. R. Ellis, K. Enqvist, D. V. Nanopoulos and F. Zwirner, *Observables in Low-Energy Superstring Models*, *Mod.Phys.Lett.* **A1** (1986) 57.
- [2] R. Barbieri and G. Giudice, *Upper Bounds on Supersymmetric Particle Masses*, *Nucl.Phys.* **B306** (1988) 63.
- [3] M. Papucci, J. T. Ruderman and A. Weiler, *Natural SUSY Endures*, *JHEP* **1209** (2012) 035, [[1110.6926](#)].
- [4] S. Fichet, *Quantified naturalness from Bayesian statistics*, *Phys. Rev.* **D86** (2012) 125029, [[1204.4940](#)].
- [5] P. Grothaus, M. Lindner and Y. Takanishi, *Naturalness of Neutralino Dark Matter*, *JHEP* **07** (2013) 094, [[1207.4434](#)].
- [6] C. Cheung, L. J. Hall, D. Pinner and J. T. Ruderman, *Prospects and Blind Spots for Neutralino Dark Matter*, *JHEP* **05** (2013) 100, [[1211.4873](#)].
- [7] T. Cohen and J. G. Wacker, *Here be Dragons: The Unexplored Continents of the CMSSM*, *JHEP* **09** (2013) 061, [[1305.2914](#)].
- [8] C. Boehm, P. S. B. Dev, A. Mazumdar and E. Pukartas, *Naturalness of Light Neutralino Dark Matter in pMSSM after LHC, XENON100 and Planck Data*, *JHEP* **06** (2013) 113, [[1303.5386](#)].
- [9] D. Barducci, A. Belyaev, A. K. M. Bharucha, W. Porod and V. Sanz, *Uncovering Natural Supersymmetry via the interplay between the LHC and Direct Dark Matter Detection*, *JHEP* **07** (2015) 066, [[1504.02472](#)].
- [10] PLANCK collaboration, P. A. R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, [1502.01589](#).
- [11] P. Gondolo and G. Gelmini, *Cosmic abundances of stable particles: Improved analysis*, *Nucl. Phys.* **B360** (1991) 145–179.
- [12] T. Cohen, M. Lisanti, A. Pierce and T. R. Slatyer, *Wino Dark Matter Under Siege*, *JCAP* **1310** (2013) 061, [[1307.4082](#)].

- [13] J. Fan and M. Reece, *In Wino Veritas? Indirect Searches Shed Light on Neutralino Dark Matter*, *JHEP* **1310** (2013) 124, [[1307.4400](#)].
- [14] A. Hryczuk, I. Cholis, R. Iengo, M. Tavakoli and P. Ullio, *Indirect Detection Analysis: Wino Dark Matter Case Study*, [1401.6212](#).
- [15] N. Arkani-Hamed, A. Delgado and G. Giudice, *The Well-tempered neutralino*, *Nucl.Phys.* **B741** (2006) 108–130, [[hep-ph/0601041](#)].
- [16] P. Ciafaloni and A. Strumia, *Naturalness upper bounds on gauge mediated soft terms*, *Nucl.Phys.* **B494** (1997) 41–53, [[hep-ph/9611204](#)].
- [17] J. A. Casas, J. M. Moreno, S. Robles, K. Rolbiecki and B. Zaldivar, *What is a Natural SUSY scenario?*, *JHEP* **06** (2015) 070, [[1407.6966](#)].
- [18] M. E. Cabrera, J. A. Casas and R. Ruiz de Austri, *Bayesian approach and Naturalness in MSSM analyses for the LHC*, *JHEP* **03** (2009) 075, [[0812.0536](#)].
- [19] M. E. Cabrera, J. A. Casas and R. Ruiz d Austri, *MSSM Forecast for the LHC*, *JHEP* **05** (2010) 043, [[0911.4686](#)].
- [20] A. Crivellin, M. Hoferichter, M. Procura and L. C. Tunstall, *Light stops, blind spots, and isospin violation in the MSSM*, *JHEP* **07** (2015) 129, [[1503.03478](#)].
- [21] B. Allanach, *SOFTSUSY: a program for calculating supersymmetric spectra*, *Comput.Phys.Commun.* **143** (2002) 305–331, [[hep-ph/0104145](#)].
- [22] G. Blanger, F. Boudjema, A. Pukhov and A. Semenov, *micrOMEGAs4.1: two dark matter candidates*, *Comput. Phys. Commun.* **192** (2015) 322–329, [[1407.6129](#)].
- [23] G. Belanger, F. Boudjema, A. Pukhov and A. Semenov, *micrOMEGAs.3: A program for calculating dark matter observables*, *Comput. Phys. Commun.* **185** (2014) 960–985, [[1305.0237](#)].
- [24] F. Feroz and M. P. Hobson, *Multimodal nested sampling: an efficient and robust alternative to MCMC methods for astronomical data analysis*, *Mon. Not. Roy. Astron. Soc.* **384** (2008) 449–463, [[0704.3704](#)].
- [25] F. Feroz, M. P. Hobson and M. Bridges, *MultiNest: an efficient and robust Bayesian inference tool for cosmology and particle physics*, *Mon. Not. Roy. Astron. Soc.* **398** (2009) 1601–1614, [[0809.3437](#)].
- [26] F. Feroz, M. P. Hobson, E. Cameron and A. N. Pettitt, *Importance Nested Sampling and the MultiNest Algorithm*, [1306.2144](#).
- [27] LUX collaboration, D. S. Akerib et al., *Improved WIMP scattering limits from the LUX experiment*, [1512.03506](#).
- [28] LUX collaboration, A. Manalaysay, *Dark-matter results from 332 new live days of lux data*, in *Identification of Dark Matter 2016*, The University of Sheffield, Sheffield, U.K., July, 2016.
- [29] XENON1T collaboration, E. Aprile, *The XENON1T Dark Matter Search Experiment*, [1206.6288](#).
- [30] XENON collaboration, E. Aprile et al., *Physics reach of the XENON1T dark matter experiment*, *JCAP* **1604** (2016) 027, [[1512.07501](#)].
- [31] P. Cushman, C. Galbiati, D. McKinsey, H. Robertson, T. Tait et al., *Working Group Report: WIMP Dark Matter Direct Detection*, [1310.8327](#).

- [32] J. Casas, J. R. Espinosa and I. Hidalgo, *Implications for new physics from fine-tuning arguments. II. Little Higgs models*, *JHEP* **0503** (2005) 038, [[hep-ph/0502066](#)].
- [33] I. Antoniadis, C. Munoz and M. Quiros, *Dynamical supersymmetry breaking with a large internal dimension*, *Nucl. Phys.* **B397** (1993) 515–538, [[hep-ph/9211309](#)].
- [34] ATLAS collaboration, G. Aad et al., *Search for strong production of supersymmetric particles in final states with missing transverse momentum and at least three b-jets at $\sqrt{s}=8$ TeV proton-proton collisions with the ATLAS detector*, *JHEP* **10** (2014) 024, [[1407.0600](#)].